Recent Product Enhancements from Northfield Research

Anish R. Shah, CFA
Northfield Information Services
Anish@northinfo.com

Boston
Sept 13, 2011

Northfield
INFORMATION SERVICES, INC.

Overview

3 Parts:

1. Conditioning models with current market information
   – Northfield Adaptive Near Horizon Models

2. Estimation error
   – How estimation error affects optimization
   – Analytical tools added to Northfield optimizer
   – "Alpha alignment"

3. Next generation of Northfield models

Northfield
INFORMATION SERVICES, INC.

Part I: Conditioning Models with Market Information

- For more detail, see
  - Short-Term Risk from Long-Term Models
  - http://northinfo.com/documents/286.pdf

1. Do you carry an umbrella?
   - Rained yesterday but sunny for 2 weeks prior
   - Conventional model (even if updated daily) reacts slowly
   - Model's user is soaked
   - *(or in reverse - model is overly cautious after spell of rain)*

2. Do you relocate?
   - Generally sunny the past few months
   - Weatherman says tropical storm will hit
   - Model (even if updated daily) ignores information outside limited set
   - Model's user is flooded

Conditioning with Market Information (2)

- Markets are dynamic
  - Information arrives all of a sudden; effects persist

- Risk models come from a limited set of not so timely data
  - Time series of (monthly/weekly/daily) returns
  - Financial statement data (note: updating daily adds mostly noise to ratios involving price)

- The window used to build the model, e.g. years of monthly observations, makes changing directions slow
  - Even under small observation intervals, e.g. daily, the window can't be sufficiently shrunk to be responsive. The model needs data points to learn structure, e.g. correlations among factors

Northfield INFORMATION SERVICES, INC.

Conditioning with Market Information (3)

- The limited data set used to build conventional models ignores current, potentially important information
  - Volatility indices (VIX, VXN, VSTOXX, VDAX,…), option implied volatilities
  - Cross-sectional spread of returns
  - Volatility estimators using intraday high/low/open/close

- How can a model incorporate this information?

- Northfield's US Short-Term Model (1998) uses option implied volatilities

- Generalized framework (2007) works with all models
  Mathematical details in Short-Term Risk from Long-Term Models

Northfield INFORMATION SERVICES, INC.

In the Short-Term (5 Day Forecast), Current Information Beats History

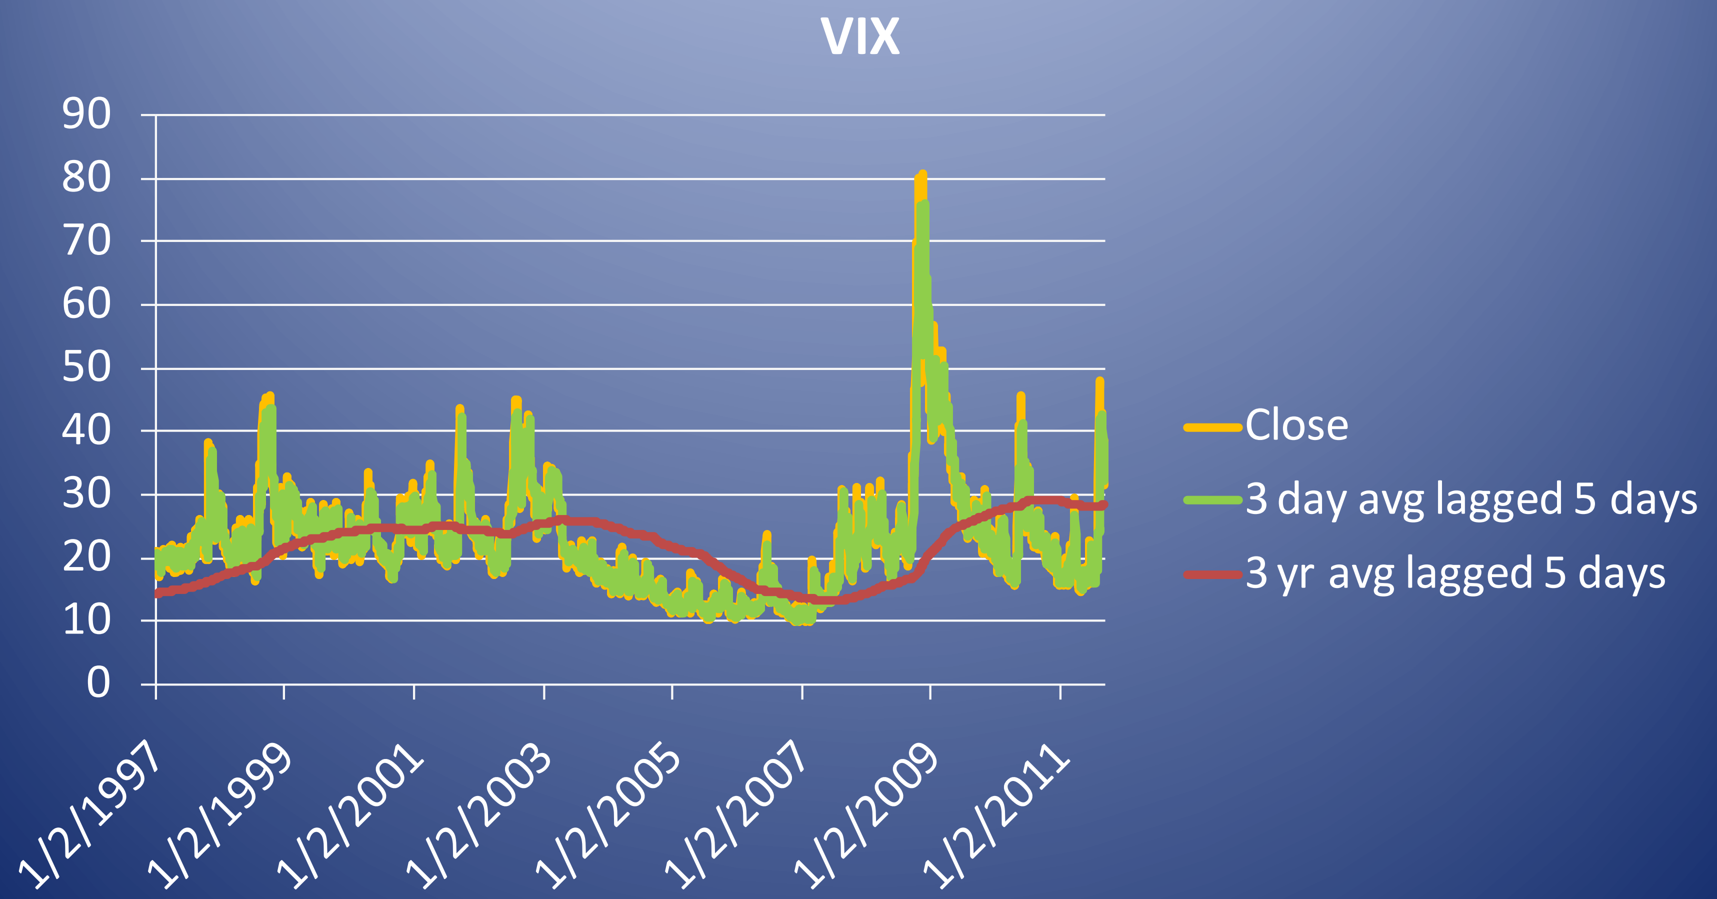

Northfield INFORMATION SERVICES, INC.

Idea in Brief

- The market information captures the current market - correlation, volatility levels (by country, sector, style, …) and more
- It comes in the form of estimates of quantities that can be predicted by the model
  - e.g. variance and cross-sectional spread of many market indices
  - exchange rate volatilities
  - interest rate volatilities
- Parameters in the risk model (e.g. variance of the European market factor, stock-specific risk levels, …) – which represent the current state of the market – are adjusted to make the model's predictions agree with the observed information
  - e.g. When VIX and other indicators say volatility has doubled, the model adapts to predict doubled volatility
  - The model instantaneously reflects current conditions

Northfield
INFORMATION SERVICES, INC.

Adaptive Near Horizon Models

- Northfield uses current market information to build daily updated, 2-week horizon versions of nearly all our models

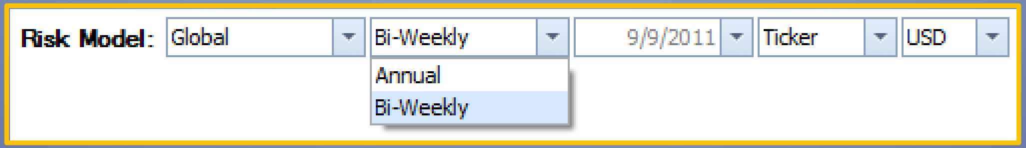

- Not for everyone! Suitable for managers with short investment horizons or leverage
  - A short-term model cares about transient effects
  - The turnover from a daily updated model can batter a long-term manager – akin to crossing the Pacific in a cigarette speedboat
  - However, any manager can use them to understand transient behavior

- A muzzled version of incorporating market information is under investigation for our long-term models

Northfield
INFORMATION SERVICES, INC.

Information Doesn't Hold for a Long Horizon (1 Yr)

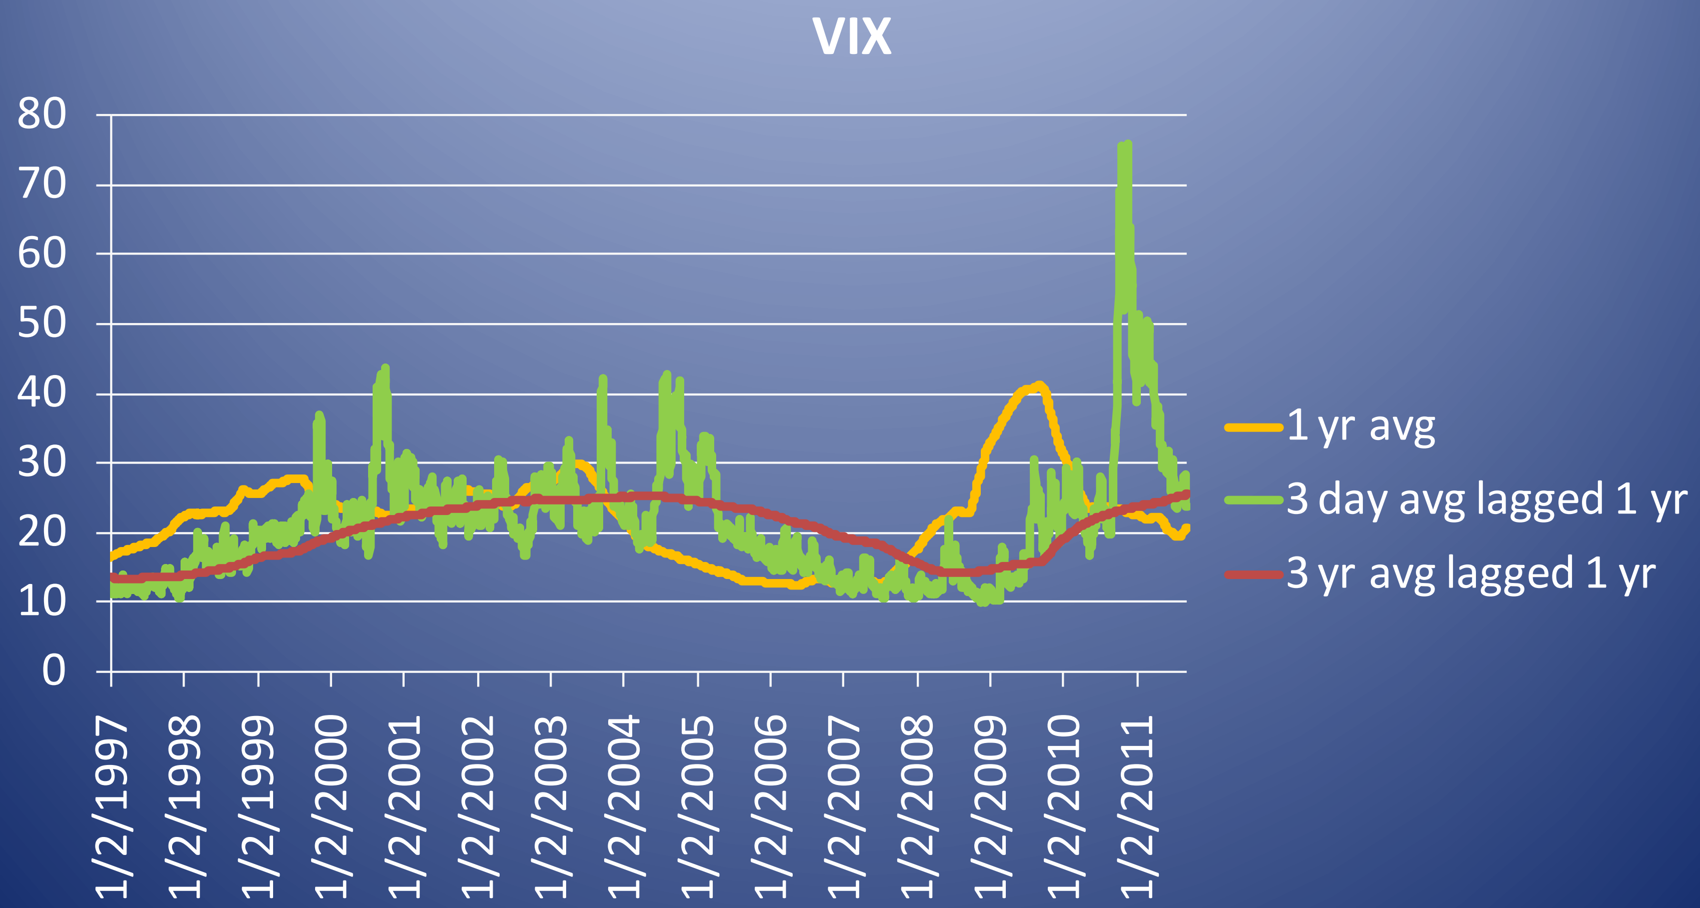

Northfield INFORMATION SERVICES, INC.

Part II – Product Enhancements to Account for Estimation Error

- For more, see:

  [Mitigating Estimation Error in Optimization]

  http://northinfo.com/documents/370.pdf

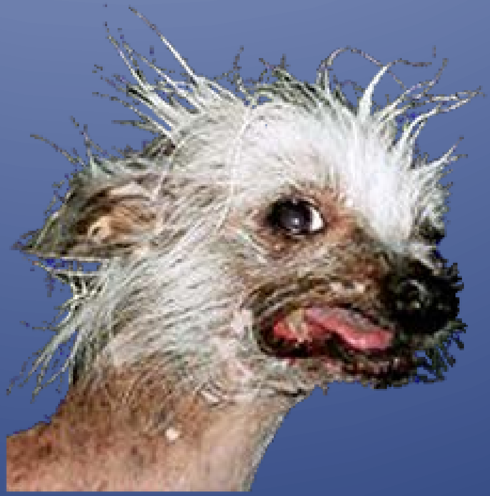

Northfield
INFORMATION SERVICES, INC.

Background in 1 Slide

- For the purpose of construction, value is measured by a single number, portfolio utility, which typically is return penalized for risk
- Optimizers find the portfolio weights that maximize this utility
- Alas, inputs into utility are only estimates
- Optimized portfolios perform worse than their forecasted utility and may look funny

Northfield
INFORMATION SERVICES, INC.

Observations

1. No single "optimal portfolio", but a large set of indistinguishably good portfolios
   - Built from forecasts, utility is an estimate (with error)
   - The error often dominates utility differences between portfolios

2. Optimal utility can be achieved by portfolios of vastly differing composition
   - An optimized portfolio's *composition* is acutely sensitive to inputs. Utility isn't
   - Nothing nefarious - the sensitivity happens because securities are alike in the characteristics determining utility (return and covariance)
   - It presents opportunity: get optimal with low transaction costs and meeting constraints
   - A consequence of the sensitivity: *it is implausible that having the same risk model as your neighbor will cause you to make the same trades*

3. Unbiased estimates are unsuitable for optimization

Northfield
INFORMATION SERVICES, INC.

Recall Definition of Unbiased Estimator

- $\hat{g}$ is an unbiased estimator of g if $E[\hat{g}] = g$
  - i.e. if I run the experiment many times, the average of my estimates ($\hat{g}_1$, $\hat{g}_2$, …) will be the true value, g
  - The error of each estimate can be huge
  - Investment managers generally have too few estimates to reach accuracy through averaging
- $\hat{g} = E[g|\text{information}]$ is a different statement, which has nothing to do with bias
  - Moreover, talking about E[g] is gibberish until the context becomes Bayesian

Unbiased Estimator: Height of Duck (Excluding Legs)

- Duck is sitting on ocean
- We are on land, eyes at sea level

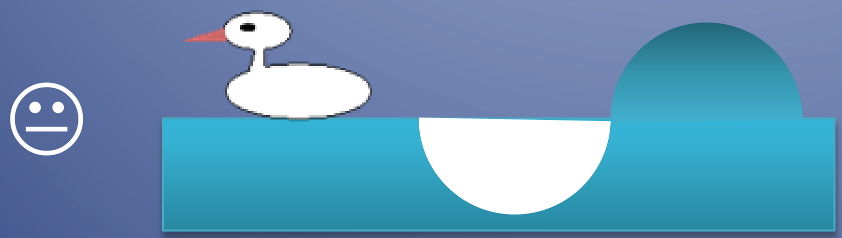

Estimated height .5 m 0 m 1 m

Average of the estimates is .5 m, the duck's true height

Add 2nd duck:

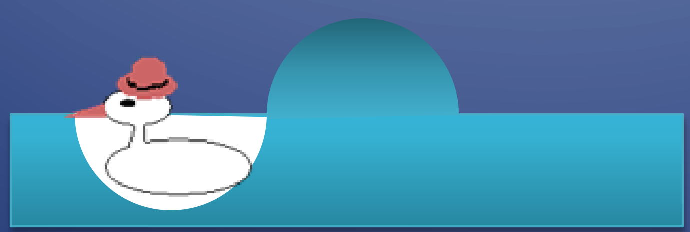

.1 m 1.1 m .6 m

Average of the estimates is .6 m, the duck's true height

Northfield
INFORMATION SERVICES, INC.

New Question: Height of Taller Duck?

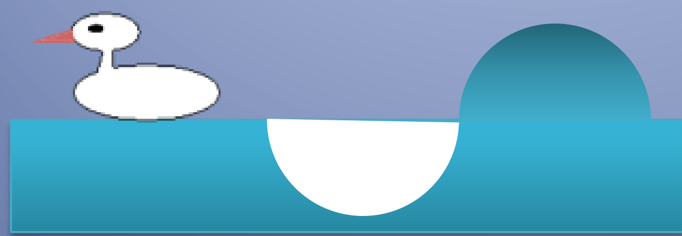

.5 m 0 m 1 m

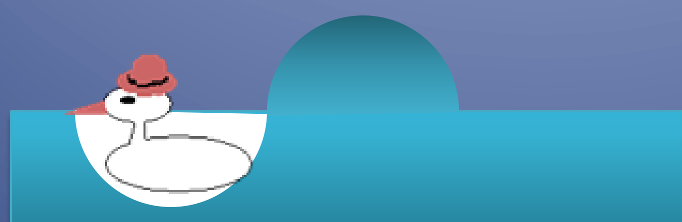

.1 m 1.1 m .6 m

Winner:

.5 m 1.1 m 1 m

Average of estimate of taller duck is 0.87 meters!

The taller duck is only .6m

The estimate is biased

Northfield
INFORMATION SERVICES, INC.

Straightforward Implications of *Unbiased* Estimates

- Markowitz utility

  $U(w,r,C) = \text{return} - \lambda \times \text{variance}$

  $= w^T r - \lambda\, w^T C w$

1) Good news: For a fixed portfolio, estimated utility is unbiased

2) Bad news: For an optimized portfolio, estimated utility on average exaggerates true utility

3) Worse news: Not only is the optimized portfolio not optimal, but its estimated utility on average exceeds the maximum achievable utility

Northfield
INFORMATION SERVICES, INC.

Real Life Effect of Optimizing with Unbiased Estimates

- What happens when utility is exaggerated?
  - Markowitz utility has 2 parts, risk and return
    - Return is on average less than estimated
    - and/or Risk is on average greater than estimated
  - Portfolio is more aggressive than preference used in the utility function
    - And (barring a fluke) not optimal at that risk tolerance
- Effect is a mathematical truth
- Easy solution: For optimization, avoid unbiased estimates

Northfield
INFORMATION SERVICES, INC.

Bayesian Inference To Limit Errors

- Classical (also called frequentist) statistics vs. Bayesian
  - Frequentist uses only observations
  - Bayesian has prior beliefs about the likelihood of events. To infer reality, Bayesian combines observations with beliefs
- Oversized hairy biped spotted in the yard
  - Frequentist yells, “Sasquatch!”
  - Bayesian, believing bigfoot sighting nearly impossible, thinks, “Mother-in-law has stopped by”

Northfield
INFORMATION SERVICES, INC.

Tool: Bayes Adjust Alphas

- Idea is described in paper by Black & Litterman (1992) but predates the Kalman Filter (1960)
  - Black, F. & Litterman, R. "Global Portfolio Optimization," Financial Analysts Journal, 1992, v48(5,Sept/Oct), 22-43

- Imagine tracking a collection of sailboats crossing the Atlantic

- The frequentist way of inferring the boats' locations
  - Receive unreliable report – "Boat X is at position Y"
  - Source sometimes overshoots, sometimes undershoots. Take report at face value and estimate boat is at position Y

- A Bayesian way
  - Know the boats' courses, hence where they intend to be today [center of prior]
  - Currents and wind affect all boats [covariance of prior]
  - Receive unreliable report of their locations [observations with error]
  - The best guess of location combines the observations and the prior belief

Northfield
INFORMATION SERVICES, INC.

Bayes Adjust Alphas (2)

- Prior on mean returns

  $m \sim N(m_0, \Sigma/\tau)$

  τ = intensity of prior

  m_0 = 0 or implied alphas of equilibrium portfolio

  Σ = covariance of benchmark relative returns

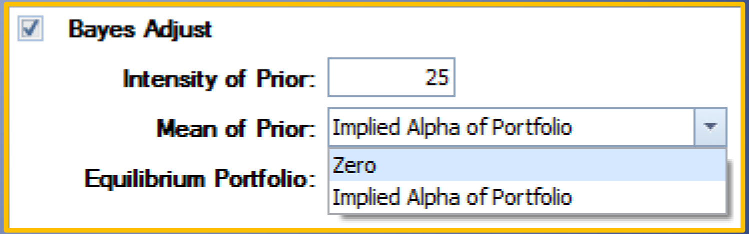

- Forecasts impart new information but are noisy observations of reality
  - $\hat{g} = m + \varepsilon \quad \varepsilon \sim N(0, \Omega)$
  - With each forecast, user provides standard deviation of the error
  - Users likely don't have opinions about the covariance of the errors, so Northfield assumes the errors are uncorrelated

- A Bayes forecast combines the prior and the error
  - Most likely returns = $m_0 + \Sigma\,[\tau\Omega + \Sigma]^{-1}\,(\hat{g} - m_0)$

Northfield INFORMATION SERVICES, INC.

Tool: Blend Covariance

- Idea comes from a series of papers by Ledoit & Wolf
  - "Improved Estimation of the Covariance Matrix of Stock Returns with an Application to Portfolio Selection," Journal of Empirical Finance, Dec v10(5):603–621
- To soften extremes, blend original covariance matrix with duller (less differentiated) versions of itself
  - Single Index
  - Constant Correlation
  - Constant Covariance

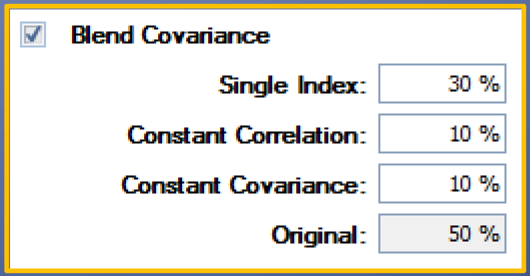

- Fortunately, each of these can be represented as a one factor model

Northfield
INFORMATION SERVICES, INC.

Blend Covariance (2)

1) Single Index

Reduce the multifactor risk model to CAPM

– $r_s = \beta_s r_m + \varepsilon_s$

– $\beta_s = cov(r_m, r_s) / var(r_m)$

– $var(\varepsilon_s) = var(r_s) - \beta_s^2 var(r_m)$

– Reference market portfolio contains all stocks in the optimization (excluding cash) weighted by cap

Northfield
INFORMATION SERVICES, INC.

Blend Covariance (3)

2) Constant Correlation

More restrictive than CAPM - all stocks have the same pairwise correlation, ρ, but different variances

– $r_s = (\rho^{½}\ \sigma_s / \sigma_f)\ r_f + \varepsilon_s$ $var(\varepsilon_s) = (1-\rho)\ \sigma_s^2$
where r_f is an artificial factor

– ρ is backed out from variance of reference portfolio
$\sigma^2_m = \sum_i (w_i \sigma_i)^2 + 2 \sum_{i<k} (w_i\ w_k\ \rho\ \sigma_i\ \sigma_k)$

→ $\rho = [\sigma^2_m - \sigma^2_{m0}] / [\sigma^2_{m1} - \sigma^2_{m0}]$
where
$\sigma^2_{m0} = \sum_i (w_i \sigma_i)^2$ = var of reference if uncorrelated
$\sigma^2_{m1} = (\sum_i w_i\ \sigma_i)^2$ = var of reference if perfectly correlated

Northfield INFORMATION SERVICES, INC.

Blend Covariance (4)

3) Constant Covariance

Least differentiated of the three – all stocks have the same pairwise correlation, ρ, and the same variance, σ^2

– $r_s = (\rho^{1/2}\ \sigma/\sigma_f)\ r_f + \varepsilon_s$ $\quad\quad\quad$ $var(\varepsilon_s) = (1-\rho)\ \sigma^2$

where r_f is an artificial factor

– σ and ρ are backed out from variance of reference portfolio

$\sigma^2_m = \sum_i (w_i\sigma)^2 + 2 \sum_{i<k} (w_i\ w_k\ \rho\ \sigma^2)$

ρ = as in constant correlation

$\rightarrow$ $\sigma^2 = \sigma^2_m / [\rho + (1 - \rho)\ \Sigma_i w_i^2]$

Northfield INFORMATION SERVICES, INC.

Sidebar: Alpha Alignment (perhaps overblown)

- What it isn't:
  1. "Having factors I use to forecast returns in the risk model eats my alpha"

     That's what risk-return tradeoff is supposed to do! Risks don't vanish via ignorance

  2. "The risk model doesn't have the factors I use to build alpha"

     Most co-movement is captured by a few (<7) factors

- What it is:
  – The factors behind the risk model and alpha forecasts are conceptually the same but defined differently, e.g. measured over 6 mos vs. 1 yr

  – Will the difference lead to situations where securities appear to have the same risk exposures but different returns, i.e. a fake free lunch?

Northfield INFORMATION SERVICES, INC.

Penalize Implied Alpha?!

- Implied alpha
  - Measures risks of positions by the alphas required to make holding them optimal
  - Analogous to measuring ecological footprint by the hectares of rainforest required to replenish resources consumed

    big implied alphas = big risks (not big returns)

    many hectares = big polluter (not an expanse of land)
  - Comes from positions and their risks. *Has nothing to do with real or forecast alphas*

Northfield INFORMATION SERVICES, INC.

Penalize Implied Alpha?! (cont.)

- A little linear algebra
  - w = (n × 1) active portfolio weights
  - C = (n x n) covariance of security returns (expressed as a factor model) $= EFE^T + \Lambda$
  - E = (n × k) factor exposures
  - F = (k × k) covariance of factors
  - Λ = (n × n) stock-specific return variances (a diagonal matrix)
  - P = (n x n) projects onto the space orthogonal to the k vectors in E $= I - E[E^TE]^{-1}E^T$ an aside: $PP = P$

- Implied alpha = scalar × Cw
  Implied alpha orthogonal to E = scalar × PCw = scalar × $P\Lambda w$

- Penalty on norm of implied alpha orthogonal to E
  = scalar × $[P\Lambda w]^T P\Lambda w$ = scalar × $w^T\Lambda P\Lambda w$

→ Results in an ad-hoc covariance adjustment (actually a stock-specific variance penalty) that has nothing to do with return forecasts

Northfield
INFORMATION SERVICES, INC.

Comments on Alpha Alignment

1. If you're going to regularize (penalize the norms of) portfolio weights via the covariance matrix, use Northfield's Blend Covariance function, based on published research by Ledoit & Wolf

2. Alpha alignment is addressed by Northfield's Bayes-Adjust Alphas function
   – Alpha forecasts on securities having similar risk are pushed toward one another

Northfield
INFORMATION SERVICES, INC.

Part III: The Next Generation of Models (in Production or Testing)

- Factor exposures inferred via Bayesian regression
  - Estimate is blend of sensitivity evidenced by data (return history) and prior belief (average exposure within peer group - stocks in same country, sector,…)
  - Blend depends on # of observations, goodness of fit, and within group spread of exposures
  - Can cover a security as soon as it appears, without any return history
- Tested Bayesian versions of both regular (least squares) regression and median regression
  - Median regression is robust to outliers
  - … but performed worse out of sample
  - Why? Exposures are dynamic, and it may be that the outliers contain information
- In models that have blind factors, the principal components are estimated using the entire universe
  - See An Introduction to Independent Component Analysis
- Investigating incorporating market information as in Adaptive Near Horizon models

Northfield
INFORMATION SERVICES, INC.

Improvements

- 2nd Generation of US Short-Term Model
  - In production since 2009
  - Bayesian regression & PCA using entire universe
  - Works uniformly better than its predecessor: more accurately captures changes in a portfolio's volatility over time, and more accurately discriminates between high and low volatility portfolios at any point in time
- US Macroeconomic Model
  - Bayesian regression & added 5 blind factors
- Single market models
  - Bayesian regression & PCA using entire universe
- Global model
  - Bayesian regression & PCA using entire universe
  - Testing a global market factor; other factors are redefined net of global market
  - New region definitions: USA/Canada, Latin America/Caribbean, Developed Europe, Emerging Europe, Middle East/Africa, Japan, Developed Asia/Asia Pacific, Emerging Asia/Asia Pacific
  - 12 added countries: Trinidad & Tobago, Bulgaria, Croatia, Serbia, Ukraine, Kenya, Lebanon, Mauritius, Nigeria, United Arab Emirates, Kazakhstan, Vietnam

Conclusion

- Research has led to two major lines of improvement in Northfield's models and optimizer:

1. A mathematical framework to incorporate current market information is the basis of the Northfield Adaptive Near Horizon models
   - respond to information instantly
   - updated daily
   - suitable for investors having short horizons or leverage

2. Bayesian techniques mitigate the effect of estimation error
   - Run-time optimizer tools: Bayes-Adjust Alphas, Blend Covariance
   - Bayesian inference to construct the next generation of Northfield models

Northfield
INFORMATION SERVICES, INC.